

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.					
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 08/09/2002		2. REPORT DATE Progress Report		3. DATES COVERED (From - To) Oct 2001 - Jul 2002	
4. TITLE AND SUBTITLE Algorithmic Issues in Network Design and in Information Access				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER N00014-98-1-0589	
				5c. PROGRAM ELEMENT NUMBER	
				5d. PROJECT NUMBER	
6. AUTHOR(S) Tardos, Eva				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Cornell University 4126A Upson Hall Ithaca, NY 14853				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research Ballston Centre Tower One 800 North Quincy Street Arlington, VA 22217-5660				10. SPONSOR/MONITOR'S ACRONYM(S) ONR	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; distribution is Unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT There has been a few main areas of efforts. (1) understanding the tradeoffs between central allocation and the cost of relying on selfish behavior in different games, including a simple model of routing on the Internet, (2) Designing algorithms and mechanisms for optimization problems for selfish users such as scheduling, path selection, and auctions, (3) design of network topology, and server placement. I will review results in each of these three areas separately.					
15. SUBJECT TERMS flows; algorithmic game theory; mechanism design; facility location; network design					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT		18. NUMBER OF PAGES
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	UU		6
19a. NAME OF RESPONSIBLE PERSON Eva Tardos					19b. TELEPHONE NUMBER (Include area code) (607) 255-0984

20030127 088

Progress report for ONR Grant N00014-98-10589

There has been a few main areas of efforts. (1) understanding the tradeoffs between central allocation and the cost of relying on selfish behavior in different games, including a simple model of routing on the Internet, (2) Designing algorithms and mechanisms for optimization problems for selfish users such as scheduling, path selection, and auctions, (3) design of network topology, and server placement. I will review results in each of these three areas separately. At the end of this report I also list two paper [6] and [4] that has been revised during this year, and is now accepted for publication.

The grant supported Tim Roughgarden (PhD student, post-doc since May 15th), Aaron Archer (PhD student), Martin Pal (PhD student).

Cost of Selfishness in Games, and Routing

In an earlier paper [13] that has just appeared in print Tardos and Roughgarden consider a routing game in networks with delay as a performance measure. More concretely, our model is defined by a graph with k source-sink pairs and traffic rates for each pair. We assume that each link has a delay that is a continuous and monotone function of the amount of flow (traffic) of the link. In the absence of regulation by some central authority, we assume that each network user routes its traffic on the minimum-latency path available to it, given the network congestion caused by the other users. In general such a "selfishly motivated" assignment of traffic to paths will not minimize the total latency. A well-known example is the Braess paradox, see Figure 1, where the edges are labeled with their latency as a function of the link congestion x . Assume that the total traffic rate is 1. In the network on the left, selfish users would split the total traffic $1/2$ - $1/2$ between the two (s, t) -path, and hence traffic on both paths would suffer a delay of 1.5. Now consider the network on the right, when a new edge with latency 0 is added. In this network all traffic would follow the paths s, v, w, t , and hence have a delay of 2. The apparent paradox is that adding an extra edge caused slower traffic for all users. In the paper [13] we consider the issue of how the "selfishly motivated" assignment of traffic fairs in minimizing the total latency, i.e., what is the cost of the lack of regulation in network performance. Our main result is that, while selfish routing can be arbitrarily more expensive than the minimum possible latency, its cost is bounded by the optimum cost of routing twice as much traffic.

Recently we made progress on two related issues. In the paper [14] we extend all the results from routing on graphs to a general class of games, nonatomic congestion games. Nonatomic congestion games have recently received attention in the game theory literature. Such games can be used to model all economic scenarios where each individual decision maker contributes only a negligible amount to the overall outcome, such as routing paths in the Internet, or many games business decisions. Equilibria in non-cooperative games are typically inefficient, as illustrated by the Prisoner's Dilemma. In [14], we quantify this inefficiency by comparing the payoffs of equilibria to the payoffs of a "best possible" outcome. We study a nonatomic version of the congestion games defined by Rosenthal, and identify games in which equilibria are *approximately optimal* in the sense that no other outcome achieves a significantly larger total payoff to the players—games in which optimization by individuals approximately optimizes the social good, in spite of the lack of coordination between players. In particular, we prove that the cost of selfish behavior is bounded by the optimal cost in the game with twice as many players.

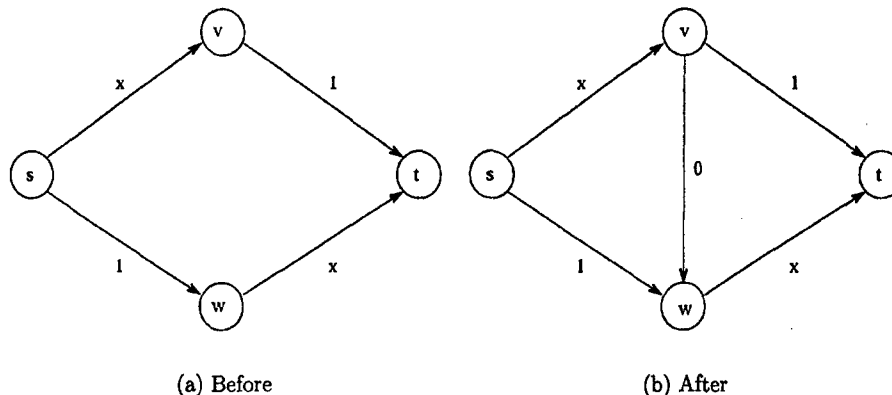


Figure 1: Braess's Paradox

In the paper [12] Roughgarden considers the cost of selfishness for limited classes of cost or delay function. He proves that for any class of convex delay functions, the underlying network topology plays no role in the determination of the cost of selfishness. Specifically, we show that the worst-case ratio between the total latency of a Nash equilibrium and of a minimum-latency routing for any multicommodity flow network is achieved by a single-commodity instance on a network of parallel links. The same result extends to the more general class of nonatomic games considered above. Roughgarden was awarded the Danny Lewin best student paper award for this paper.

In [11] Roughgarden considers the issue on how fair the optimal traffic flow is. A routing of traffic that is optimal from the viewpoint of total latency may be "unfair" in the sense that some traffic may incur far more latency in the optimal flow than in a selfishly-defined equilibrium. We prove limits on this unfairness when all traffic shares the same source and destination.

In [10] Roughgarden considers the problem of designing networks with selfish users in mind. Given a routing network with congestion dependent edge-latencies the Braess paradox shows that the deleting edges can improve the network performance. In [10] Roughgarden proves that unless $P=NP$ no polynomial time algorithm can find such improved subgraphs, even in networks where the severity of the Braess paradox grows with the network size.

Algorithms for Selfish Users

In this section we report on work in the general area of mechanism design. The goal of mechanism design is to develop protocols where selfish agents find it in their self-interest to cooperate for the desired outcome. For example, in the above routing game, if we were allowed to charge each agent with a fee for using edges, and if we assume that the goal of a selfish agent is to minimize the sum of his delay and the cost, then by introducing appropriate costs one can guarantee that each agent would find it in their interest to follow the route assigned to him in a globally optimal routing.

We have been working on such mechanism design questions related to scheduling, routing, and

auctions. In each game the goal of the mechanism designer is to solve a natural combinatorial optimization problem. However, in addition to aiming to optimize (or approximately optimize) this goal, the algorithm also has to guarantee that it is never in any agents interest to deviate, e.g., by reporting false data. We call such an algorithm *truthful*. The main positive tool in mechanisms design is the celebrated Vickrey-Clark-Grove mechanisms (or VCG mechanisms, for short).

Unfortunately, the VCG mechanisms have a number of shortcomings. They can only handle objective functions of a certain form, namely maximizing the average happiness of all agents (traditionally referred to as the *social welfare*, or *utilitarian* objective). There are many other types of goals that naturally arise in practical settings. The algorithm assumes that the above utilitarian objective function can be optimized efficiently, while in many cases this optimization problem is NP-hard. Also, the mechanisms uses side-payments to the agents to induce the desired behavior, and at times these payments can be very high. The general goal of our work in this area is to develop truthful mechanisms for many important optimization problems that alleviate the above problems with a VCG type mechanism. An important tool that we will employ is approximation algorithms. Most of the optimization problems that natural mechanisms aim to optimize are NP-hard. Over the last 15-20 years there has been a large amount of work on approximation algorithms for a huge array of hard optimization problems. However, so far there are only very few examples when such approximation algorithms turned out to be useful for truthful mechanisms.

In the paper [1] we consider the a simple scheduling problem: scheduling on related parallel machines. Here each job j has a (publicly known) processing time p_j . The machines are each owned by separate agents, and the agent's secret data is naturally expressed by a single positive real number, its speed s : if a job of processing time p is assigned to an agent that has speed s than the agent can process the job in sp time. The problem of minimizing makespan for m related machines remains NP-complete, however, there are a number of simple approximation algorithms for the problem, and there is also a polynomial time approximation scheme. We show that these approximation algorithms cannot be made truthful by introducing side-payments, and develop a truthful 3-approximation algorithm for the problem.

In the paper [2] we consider consider the shortest path problem in this game theoretic context. Assume that the costs c_e of the edges $e \in E$ are not directly available to the algorithm. Instead, there are several economic agents who each "own" some of the edges. The agent who owns edge e suffers a cost of c_e if edge e is selected, but he may lie when reporting this cost to the algorithm. In the context of the shortest path problem, maximizing the "social welfare" coincides with the natural objective of finding a path of minimum cost. However, truthfulness is achieved in the VCG mechanism through using side payments to the players, analogous to the arc costs suggested in the routing example in the beginning of this section. In general, side payments are necessary to achieve a desirable outcome via a truthful scheme. However, the side-payments used by the VCG mechanisms are often excessive, and hence it leads to a truthful, and socially desirable, but financially expensive solution. This is the case, for example, for the shortest path problem, where the VCG mechanism can be forced to pay a huge amount, compared not only to the actual cost of the path, but also compared to the cost of the cheapest path that is completely disjoint from it. In the paper [2] we show roughly, that all reasonable path selection mechanisms can be forced to over-pay just as badly as VCG.

In [3] we develop a general technique for turning certain approximation algorithms into truthful mechanisms, and apply it for a special case of the auction problem. Randomized rounding is a

commonly used technique for design of approximation algorithms. We show how to make a rounding procedure incentive compatible, and hence give a truthful mechanism for combinatorial auctions with single parameter agents (e.g., single minded bidders) that approximately maximizes the social value of the auction. We consider combinatorial auctions where multiple copies of many different items are on sale, and each bidder i desires a subset S_i . Given a set of bids, the problem of finding the allocation of items that maximizes total valuation is the well-known set-packing problem. This problem is NP-hard, but for the case of items with many identical copies the optimum can be approximated very well. To turn this approximation algorithm into a truthful auction mechanism we overcome two problems: we show how to make this algorithm monotone, and give a method to compute the appropriate payments efficiently.

Network Design

During this year we have worked on network design problems whose goal is to design networks that can carry traffic efficiently. Consider the following general network design problem. We are given a graph G with costs c_e on the edges, and k commodities, so that we need to route some given amount of flow between the source-sink pairs of these commodities. We consider a very simple model of the natural economy of scale: the cost for x unit of flow on edge e will be $\min(x, M)c_e$, and the goal of the problem is to find a multicommodity flow of minimum cost. Note that these costs are concave (modeling a natural economy of scale), and as a result this optimization problem is NP-complete. The best approximation algorithm known for this problem was an $O(\log k \log \log k)$ approximation due to Bartal. In the special case when all commodities share a source, constant factor approximation has been obtained recently, even for more general concave separable objective functions. In the paper [7] we give a constant factor approximation for the above multicommodity network design problem.

In [5] we consider the problem of finding a Steiner tree that contains at least k of a set of given points (called the k -Steiner tree problem). Naveen Garg gave a 5-approximation algorithm for this problem in FOCS'96. Here we show that Garg's algorithm can be explained simply with the Lagrangean relaxation framework introduced by Jain and Vazirani for the k -median problem.

In [8] we consider a simple directed network design problem from a computational perspective. We consider the problem of finding a minimum cost strongly connected subgraph of a given directed graph. The best approximation algorithm known for this problem is a 2-approximation that combines the minimum cost in-branching and out-branching in the graph. However, this algorithm is experimentally quite bad (that is, its performance is often close to the provable bound of 2). In [8] we develop primal-dual-based algorithms for the problem, that both have a constant worst case guarantee (though slightly worse than the guarantee of the above branching based algorithm) and also has good performance in practice.

In [9] we used local search to develop an approximation algorithm for the facility location problem with hard capacities. Korupolu, Plaxton, and Rajaraman showed that a simple local search heuristic proposed by Kuehn and Hamburger in 1963 provides a constant factor approximation for facility location problems with uniform hard capacities. Chudak and Williamson improved on the analysis. However, the Kuehn and Hamburger heuristic is not strong enough to handle non-uniform capacities. In [9] we design a more powerful version of the local search heuristic, and prove that it yields a $9 + \epsilon$ approximation algorithm for the problem with non-uniform, hard capacities.

References

- [1] Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, pages 482–491, 2001. <http://www.orie.cornell.edu/archer/Research/oneparamFOCSproc.ps>
- [2] Aaron Archer and Éva Tardos. Frugal path mechanisms. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 991–999, 2002. <http://www.orie.cornell.edu/archer/Research/fpSODAprc.ps>
- [3] Aaron Archer, Christos Papadimitriou, Kunal Talwar, Eva Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents.
- [4] M. Charikar, S. Guha, E. Tardos and D. Shmoys. A constant-factor approximation algorithm for the k-median problem To appear in the STOC 2000 special issue of the *Journal of Computer and System Sciences*. <http://www.cs.cornell.edu/home/eva/median.ps>.
- [5] F. A. Chudak, T. Roughgarden, D. P. Williamson. Approximate k-MSTs and k-Steiner trees via the primal-dual method and Lagrangean relaxation. in the Workshop on Integer Programming and Combinatorial Optimization (IPCO'01) 2001 pages 60–70. http://www.cs.cornell.edu/timr/papers/kmst_full.ps.
- [6] J. Kleinberg and E. Tardos. Approximation Algorithms for Classification Problems with Pair-wise Relationships: Metric Partitioning and Markov Random Fields To appear in the Journal of the ACM. <http://www.cs.cornell.edu/Info/People/eva/papers/vision.ps>
- [7] Amit Kumar, Anupam Gupta, and Tim Roughgarden. A Constant-Factor Approximation Algorithm for the Multicommodity Rent-or-Buy Problem, To appear in the Proceedings of the Annual IEEE Symposium on the Foundation of Computer Science (FOCS), 2002.
- [8] Vardges Melkonian and Eva Tardos. Primal-Dual-based algorithms for a directed network design problem. Submitted for journal publication.
- [9] M. Pal, É. Tardos, and T. Wexler. Facility location with hard capacities, in the Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science, 2001. http://www.cs.cornell.edu/People/mpal/papers/fl_hard.ps.
- [10] T. Roughgarden. Designing networks for selfish users is hard. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science*, pages 472–481, 2001. http://www.cs.cornell.edu/timr/papers/nd_hard_full.ps
- [11] T. Roughgarden. How unfair is optimal routing? In *Proceedings of the Annual Symposium on Discrete Algorithms*, 2002, pages 203–204. <http://www.cs.cornell.edu/timr/papers/unfair.ps>
- [12] T. Roughgarden. The price of anarchy is independent of the network topology. In *Proceedings of the 34th Annual ACM Symposium on the Theory of Computing*, 2002, pages 428–437. http://www.cs.cornell.edu/timr/papers/indep_full.ps.
- [13] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM* 2002, pages 236–259. <http://www.cs.cornell.edu/timr/papers/routing.ps>.

- [14] T. Roughgarden and É. Tardos. Bounding the Inefficiency of Equilibria in Nonatomic Congestion Games Submitted for journal publication.
<<http://www.cs.cornell.edu/timr/papers/ncg.ps>>

UNCLASSIFIED

[This page is intentionally left blank.]

UNCLASSIFIED